

AD 606530

GAME THEORY

Herman Kahn
Irwin Kahn

✓
P-1166

July 30, 1957

Approved for DTS release

68-10

BY	1	CF	1	5
REPRODUCTION				\$. 7.00
MICROPHONE				\$. 8.75



The RAND Corporation

1700 MAIN ST • SANTA MONICA • CALIFORNIA

Copyright 1958
The RAND Corporation

PREFACE

This is a draft of a report which is being circulated for information and comment. We hope to make it a chapter of a book titled Military Planning In An Uncertain World, and would appreciate any comments, criticism, ideas, and examples that readers may have. This draft began as a transcript of an informal talk and, despite some rewriting, it probably still suffers (like many such talks) from being "fashionable." We are aware that it has a number of other weaknesses and assume there are still others of which we are not aware. We hope to give it a thoughtful and leisurely review but are deferring this until we get some outside criticism.

A table of contents is given on the next page to show the relation of this chapter to the rest of the book. The chapter may not be quite self-contained as a paper, as it occasionally refers to other chapters; but we trust this will be understood or overlooked.

A more complete introduction and list of acknowledgements are given in RM-1829-1.

1 2

**CLEARINGHOUSE FOR FEDERAL SCIENTIFIC AND TECHNICAL INFORMATION CFSTI
DOCUMENT MANAGEMENT BRANCH 410.11**

LIMITATIONS IN REPRODUCTION QUALITY

ACCESSION #

AD 606 530

- ☒ 1. WE REGRET THAT LEGIBILITY OF THIS DOCUMENT IS IN PART UNSATISFACTORY. REPRODUCTION HAS BEEN MADE FROM BEST AVAILABLE COPY.
- ☐ 2. A PORTION OF THE ORIGINAL DOCUMENT CONTAINS FINE DETAIL WHICH MAY MAKE READING OF PHOTOCOPY DIFFICULT.
- ☐ 3. THE ORIGINAL DOCUMENT CONTAINS COLOR, BUT DISTRIBUTION COPIES ARE AVAILABLE IN BLACK-AND-WHITE REPRODUCTION ONLY.
- ☐ 4. THE INITIAL DISTRIBUTION COPIES CONTAIN COLOR WHICH WILL BE SHOWN IN BLACK-AND-WHITE WHEN IT IS NECESSARY TO REPRINT.
- ☐ 5. LIMITED SUPPLY ON HAND: WHEN EXHAUSTED, DOCUMENT WILL BE AVAILABLE IN MICROFICHE ONLY.
- ☐ 6. LIMITED SUPPLY ON HAND: WHEN EXHAUSTED DOCUMENT WILL NOT BE AVAILABLE.
- ☐ 7. DOCUMENT IS AVAILABLE IN MICROFICHE ONLY.
- ☐ 8. DOCUMENT AVAILABLE ON LOAN FROM CFSTI (TT DOCUMENTS ONLY).
- ☐ 9.

PROCESSOR: *2 1*

MILITARY PLANNING IN AN UNCERTAIN WORLD

CONTENTS OF BOOK

- I. Techniques of Systems Analysis¹
 - 1. Designing the Offense
 - 2. Probabilistic Considerations
 - 3. Designing the Defense
 - 4. The Two-Sided War
 - 5. Evaluation and Criticism
- II. Techniques of Operations Research
 - 6. Flyaway Kits--An Application of and Introduction to Chapters 7 and 8
 - 7. Elementary Economics and Programming
 - 8. Probability and Statistics
 - 9. Monte Carlo²
 - 10. Game Theory
 - 11. War Gaming³
- III. Philosophical and Methodological Comments
 - 12. Ten Common Pitfalls⁴
 - 13. Nine Helpful Hints
 - 14. Miscellaneous Comments

1. Has already appeared as RM-1829-1
2. Has already appeared as P-1165
3. Has already appeared as P-1167
4. Has already appeared as RM-1937

GAME THEORY

Game Theory is not only the study of games, per se, but is more generally the study of any conflict situation.¹ The ordinary parlor game or athletic contest is a simple and clear-cut example of a situation where people have conflicts of interest; much of the terminology is drawn from these two fields. In principle, however, we will be talking about any situation where there are two or more participants who do not have identical objectives. Thus, game situations are to be differentiated from what economists call the Robinson Crusoe Economy where a single person or monolithic group has control of all of the decisions. Chapter 7 on economics indicated some of the relevant considerations in this case. One of the major results of Game Theory is that it shows, clearly and conclusively, that such simple considerations are not sufficient to handle conflict situations--having more than one optimizing player introduces new concepts.

In Game Theory, even more than in most of the other subjects discussed in Part Two, the interest is not so much in the numerical results that can be obtained by studying specific games, but in the intellectual content of the subject. The subject matter of Game Theory is usually a highly idealized abstraction of real life. Therefore, most of the games that have been studied do not have (numerically) important normative or predictive aspects.²

¹Actually we are going to omit much of the formal theory and spend most of our time discussing some typical games. There already exists a delightful little book called "The Compleat Strategyst" by John Williams, which gives an entertaining and elementary account of some of the formal theory with many examples.

²This statement is not meant to deny that there are important numerical applications, but only implies that they do not occupy a central role. On the whole the applications to real games are dwarfed by the many insights that the theory provides.

We can, however, gain insight into the nature of conflict situations by looking at some of the theorems and then illustrating these theorems in examples. Hopefully, the understanding thus gained will guide our intuition even where the known theorems do not directly apply. It will also be as interesting to see what cannot now be done by mathematics as what can, so that our ambitions will be curbed. We hope that the insight and understanding that we obtain will train our judgment and improve our vocabulary; generally it will not give precise rules for specific realistic situations.

The technical vocabulary of Game Theory in itself is valuable. It is rich and suggestive without being ambiguous. It therefore provides a very useful tool. Where it can be used, it tends to be superior to the competing vocabulary of the Social Scientist. In the Social Sciences, partly because there is no mathematical discipline, and even more because no one feels compelled to use the results and terminology of anybody else's papers, the meaning of the terms is often not quite clear or generally accepted.³ One of the major tasks which we hope to accomplish in this section is to explain and use this vocabulary.

³ Some of our nastier friends in Social Science have told us one of the reasons for this is that, unlike the mathematics field, there is no reason to refer to anyone else's paper. This is probably too harsh a way to say it. Most social science papers have to be expository, or at least semi-expository (with a resultant inordinate increase in the total volume of words). This makes it almost physically impossible to read them all seriously. In the mathematics field one does not publish papers unless he has results (theorems) which he thinks are important and new. Except for review articles, the papers are almost bare of general exposition. Since future papers of others usually should use and always must refer to these theorems, the authors have to read the papers and are automatically introduced to the vocabulary. It is amazing and indicative how often the definitions (and even the notation) of pioneering papers are used by most later writers even when the pioneering papers were somewhat pedestrian.

Matching Pennies

About the simplest game one can think of is the ordinary matching game where I^h choose between heads and tails and my opponent makes a similar choice.^h If we make the same choice, I^h win; if not, I^h lose. One can describe a game like this by showing two arrays of numbers which are called the pay-off matrices. There are two such matrices--one for me and one for my opponent as shown below.

		his choice	
		H	T
my choice	H	+1	-1
	T	-1	+1

My Payoff Matrix

		his choice	
		H	T
my choice	H	-1	+1
	T	+1	-1

His Payoff Matrix

Chart 1

The payoff matrices show everything that can happen. If, for example, I choose heads and he chooses heads, I get a dollar, and he loses one. If I choose heads and he chooses tails, I lose a dollar and he gets one, and so on. This is an example of the so-called zero-sum two-person game. It is clear that in this case it is not necessary to show a payoff matrix for both me and my opponent since one matrix is the negative of the other; everything I win he loses and vice versa. Whenever a game has this last property it is known as a "zero-sum" game because of the obvious property that the sum of the winnings of all the players is zero.

^hOne of the authors finds this game dull. ^{hA}

^{hA}One of the authors finds the other author dull.

In the future when only one matrix is shown it should be understood that the matrix of the other player is the negative of the displayed matrix. The player whose payoff matrix is shown will be called the maximizing player and his choices will always be used to label the rows. The player who gets the negative of the amount shown is known as the minimizing player and his choices label the columns.

The zero-sum two-person game is especially simple to treat. The other situation, where there are more than two participants, or where the game is not zero-sum so that the matrices of the two opponents are essentially different, is much harder. There are, in fact, often real unanswered questions in the simplest examples of these other cases. We shall discuss some examples of these more complicated games in the second half of this chapter.

Let us go back to the game of heads and tails. We will call a decision to choose heads or tails a strategy (more accurately a pure strategy). This may seem a slightly curious use of the word, but in fact any complete and consistent set of choices that is available to a player will be called a strategy. Now, if I were going to play this matching game, it is clear that I wouldn't want my opponent to know whether I was going to pick heads or tails. The usual way to do this is just not to let him know which I have done until after he has also made his choice. If, however, I have to play the game very often, I may fall into a pattern and telegraph my choice, particularly if my opponent is a skilled psychologist. For example, if one plays this game with a child, one can predict pretty well, at least in some cases, many of the child's choices. In order to avoid any possibility of this, I may not actually make the choice myself, but use a random device to choose equally between heads and tails. The simplest such random device is to toss a coin in the air.

Another way to accomplish the same thing is to draw random numbers from a table. Every time the random number is less than half, one might pick heads, and every time it is greater, one would pick tails.

It is irrelevant what the random device is. The only thing of importance is that the probability assigned to each choice be $1/2$. This mixing of choices with probability $1/2-1/2$ (or with any probability) is also called a strategy (more accurately a mixed strategy). The reader will notice that we don't care if our opponent knows the mechanism by which we make our choice or the probabilities we use.

We have implicitly assumed in our discussion that we should use a $1/2-1/2$ mixed strategy. Why is this so? In the next chart we have shown two heads-tails games, one played with the usual $1/2-1/2$ probability and one played with probabilities $1/4-3/4$, where we put down the probabilities we have assigned to each strategy, and have put question marks for my opponent's probabilities.

strategy his choice		strategy my choice			
		H	T		
strategy	H	1/2	H	+1	-1
	T	1/2	T	-1	+1
expected value of my winnings		0	0		

strategy his choice		strategy my choice			
		H	T		
strategy	H	1/4	H	+1	-1
	T	3/4	T	-1	+1
expected value of my winnings		-1/2	+1/2		

Chart 2

On the bottom we have put the expected value of my winnings for each of the particular choices available to my opponent. (The reader should know by now how the expected values are calculated. One simply multiplies every payoff by its probability and adds them up.)

When I use the classical $1/2-1/2$ mixture, the expected values of both columns are zero and it does not make any difference which column my opponent picks. By playing the $1/2-1/2$ strategy I have left no loopholes where my opponent can get an advantage. So long as I am playing this strategy, he can play anything he wants and I will not lose. Also, no matter how stupid he is, I will not get any advantage from his stupidity. None of these remarks are true if I use the $1/4-3/4$ strategy. In that case, he can pick heads and win 50 cents a play on the average.

The contrast is typical. When I picked the $1/2-1/2$ mixed strategy, I probably implicitly assumed that the enemy was bright and would take advantage of any mistakes I made. By trying to prevent him from taking advantage of any mistakes I might make, I have lost any ability I might have had to take advantage of his mistakes. In more general cases the situation is often not this drastic; that is, if we try to protect ourselves against the possibility that the enemy is smart we ordinarily lose some in our ability to take advantage of his stupidity; we don't typically lose the ability completely.

Modified Matching Game

Ordinary matching is a pretty dull game. Let us change it a little. Let us introduce the rule that if we match him with two heads he has to give us \$2.00 while if we match him with two tails we break even. However, as before if we fail to match, we pay \$1.00. What would happen now if we tried our $1/2-1/2$ strategy. Well, Chart 3 shows it.

my choice strategy		his choice strategy		
		H	T	
1/2	H	2	-1	
1/2	T	-1	0	
column average		+.50	-.50	

Chart 3

The expected value to us, if the enemy plays the first column is +50 cents and the expected value if he plays the second column is -50 cents. It would not take very long, if we played this game and used the 1/2-1/2 strategy, for the enemy to catch on to what was happening and always play the second column. Or if he wanted to be a little deceptive he would at least play column two much more frequently than column one. In effect, he has found a loop hole in our system and will probably take advantage of it. In order to prevent this, we must change our strategy. It is fairly clear how we must change it. Chart 4 shows a graph of what happens as a function of our probabilities.




Chart 4

The heavy bent line shows the worst that the opponent can do to us. It is the line of minimum income. If we play heads less than $1/4$ of the time and he knows it, then he will also play heads. If we play heads more than $1/4$ of the time and he knows it, he will switch to tails. If we play heads exactly $1/4$ of the time, then we don't care what the opponent does and we don't care if he knows it. This $1/4$ point is the highest point on the heavy line so it is the best point for us in the sense that this is the highest average income we can guarantee ourselves. It is called the maximum of the minimums (abbreviated max-min). If we play this point then no matter how the enemy plays, one column or the other, he will get, on the average, 25 cents and no more.

We are not delirious about this. After all, we are losing, on the average, 25 cents every play of the game. However, this is the most we will lose. We no longer have a loophole, so to speak, through which the enemy can get more than that amount.

It is clear from the chart that there is no way of playing this game against a good opponent which will lose us on the average less than 25 cents. What about a bad opponent? Well, if we play our "optimum" strategy as before, it makes no difference how badly our opponent plays. However, if our opponent plays foolishly, we might switch our strategy to take advantage of it. For example, assume that he thinks that we will play the $1/2$ - $1/2$ strategy and therefore decides in turn to play the second column. He then thinks that he will win, on the average, 50 cents. Actually we could then play the second row, and he would not win anything. In other words, he should also mix his strategy between the two columns so that we cannot predict what he will do. He must pick the correct mixture too (as we shall see it happens to be the same as ours, $1/4$ and $3/4$). If he happens to play the old $1/2$ - $1/2$ strategy,

the result is as follows.

strategy his choice		strategy my choice		Expected Value of my winnings
		H	T	
?	H	2	-1	.50
?	T	-1	0	-.50

Chart 5

By playing the first row we can win 50 cents instead of losing 25 cents. But if we play the optimal $1/4-3/4$, we still lose 25 cents because we are no longer in a position to take advantage of our opponent's mistake. If we are both foolish and both play the $1/2-1/2$ strategy, then the game comes out even instead of being 25 cents against us. In these circumstances, it might be wise to continue playing the $1/2-1/2$ strategy. If we try to be clever and play a strategy in which heads are more frequent than tails, we may tip the opponent off to the fact that he should start looking around for a better strategy. If he does, he may find the optimal $1/4-3/4$ strategy and start to win.

Chart 6 shows how our opponent will look at the situation if he thinks we are good players. The bent line shows the best that we can do against any mixed strategy of his. It is the line of our maximums. If he thinks we are good players, he will assume that we will always try to play the strategy that corresponds to this line. He will, therefore, pick a strategy that puts us on the minimum point of this line of maximums. This corresponds to his playing a $1/4-3/4$ mixture of heads and tails.




Chart 6

In other words, if our opponent plays heads less than $1/4$, then we should play tails; if he plays heads more than $1/4$ of the time, we should also play heads. If he plays the $1/4$ - $3/4$ point, he doesn't care what we play. This point is the minimum of the maximums (abbreviated min-max). At this point we again lose an average of 25 cents per play. While it is a coincidence that our previous strategy associated with our max-min point is the same as his strategy associated with his min-max point, it is no coincidence that the average payoffs are the same. If we use mixed strategies in a two-person zero-sum game then the payoff associated with max-min always equals the payoff associated with min-max. The fact that there exists such an optimal strategy in this sense for both players is the fundamental theorem of two-person zero-sum games.

Such a strategy for each player is usually unique, though it is easy to construct examples where it is not. In any case, however, uniqueness holds for the expected values of the outcome (payoff).

An optimal strategy also has the advantage that it need not be kept secret (allowing that the opponent is able and willing to compute his own optimal strategy). A particular choice, however, must be secret. If a player's choice is known to the opponent, the opponent can gain an advantage by having this knowledge. In fact, the different game in which I must choose and make known my choice before he makes his choice is called the minorant game. In that game, I will simply win the maximum of the row minimums. In the majorant game, still another game, where he must choose first, I can gain the minimum of the column maximums. Chart 7(a) illustrates that this amount is more than the return from the minorant game.

Using a pure strategy is a particular mixed strategy, but in the original game, a pure strategy cannot be optimal. Using the optimal mixed strategy, it doesn't matter who knows the recipe for the mix. We are at a point called a saddlepoint, where the payoff increases if the minimizing player deviates and decreases if the maximizing player does. Sometimes a saddlepoint occurs at an element of the matrix, simply whenever the minimum of a row is also the maximum of a column. This happens when the game is modified as in Chart 7(b). I will always pick heads and my opponent tails. The game costs one dollar each time I play, and we should probably arrange for an annuity to my opponent instead.

his choice my choice				row minimum
		H	T	
H	2	-1	-1	
	T	-1	0	
column maximum		2	0	

Using pure strategies

max-min = -1

min-max = 0

-a-

A Game Without a Saddlepoint

his choice my choice				row minimum
		H	T	
H	4	-1	-1	
	T	-1	-2	
column maximum		4	-1	

Using pure strategies

max-min = -1

min-max = -1

-b-

A Game With a Saddlepoint

Chart 7

A Game of Ruin

We have not finished with this matching game. Let us change the situation a little and say that I don't want to play to win the most I can per play, but rather that I wish to ruin my opponent. Assume he has some fixed fortune, say \$2.00, and that I likewise have \$2.00. We are going to play until one or the other of us is bankrupt. At first sight the choice of game might seem somewhat unfortunate from my point of view. After all, what we previously called "optimum" play is somewhat against me. Actually, this game is a disaster; I have no chance at all of bankrupting a smart opponent. Let us see why this is so.

		his choice	
		H	T
my choice	H	2	-1
	T	-1	0

Chart 8

If we both play the old optimum $1/4-3/4$ strategy I will find that, on the average, the enemy would bankrupt me slightly more than four times out of five and I would bankrupt him slightly less than one time out of five.⁵ While this is poor it is not completely devastating. However, in this game of bankruptcy, the enemy has a much better strategy. He can simply play the second column most of the time. As long as he is playing this column, I can never force him to lose any money. If I play the first row, I will lose a dollar. If I play the second row, I will break even. Therefore, if I know he is going to play the second column most of the time, I should play the second row most of the time. As long as we are both doing this we will break even. The trouble is that once in a great while, he will suddenly shift his strategy and play the first column. Now if I happen to shift my strategy at the same time, I will win \$2.00, but if I don't, I will lose a

⁵There is a probability $1/16$ of gaining 2, $9/16$ of breaking even, and $6/16$ of losing 1. It is then easy to verify that the probabilities of going bankrupt with fortunes (1,3), (2,2), and (3,1) are .96, .82, and .70 respectively.

dollar. Furthermore, if I happen to make a mistake and shift when he does not shift; that is, if I play row one when he is still playing column two, I will also lose a dollar. In other words, I can't afford to shift unless I know the enemy is shifting and I can't afford not to shift if the enemy does shift. This is a completely intolerable situation for me and he will undoubtedly succeed in bankrupting me. Thus, the game which was somewhat unfair when played in a friendly spirit is catastrophic when the objective is no longer to maximize the winnings per play but to maximize a longer term objective, the probability of bankrupting the opponent.

In all the games which follow we will assume that our object is to maximize the average amount we can win per play and not to bankrupt our opponent, or some other objective. This is an important caveat in practice.⁶

⁶It is possible to make some (approximate) general statements about how to play games of ruin as opposed to games in which one is trying to maximize the average winnings per play.

The following is well known in mathematical literature. Consider a random walk where one wishes to calculate the probability of reaching a boundary at a point B before reaching a boundary at a point A (A corresponds to ruin, B to success). The probability of doing this is determined by the following integral equation:

$$P(x) = \int_A^B k(x,y)P(y)dy + \int_B^{\infty} k(x,y)dy \quad (1)$$

where $P(x)$ is the probability one wishes to know and $k(x,y)$ is the probability of jumping from x to y .

The equation simply states that the probability of eventually getting to a position greater than B from any point x is equal to the probability of going there immediately plus the probability of going to some other point between A and B and then getting greater than B eventually. One can expand $P(y)$ in the first integrand in a Taylor series

$$P(y) = P(x) + P'(x)(y-x) + \frac{1}{2}P''(x)(y-x)^2 + \dots \quad (2)$$

around the point x and after making some further reasonable approximations obtain the differential equation:

(footnote continued on next page)

$$\frac{1}{2}q(x)\frac{d^2P}{dx^2} + m(x)\frac{dP}{dx} = 0 \quad \begin{matrix} P(A) = 0 \\ P(B) = 1 \end{matrix}$$

where

$$\begin{aligned} m(x) &= \int_{-\infty}^{\infty} (y-x)k(x,y)dy \\ q(x) &= \int_{-\infty}^{\infty} (y-x)^2 k(x,y)dy \end{aligned} \quad (3)$$

and we are assuming that

$$\int_{-\infty}^{\infty} k(x,y)dy = \begin{cases} 1 & \text{if } x \geq B \\ 0 & \text{if } x \leq A \end{cases}$$

If m and q happen to be constant the solution of equation (2) is

$$P(x) = \frac{e^{-2mx/q} - e^{-2mA/q}}{e^{-2mB/q} - e^{-2mA/q}} \quad (4)$$

The important thing to notice is that the probability of success depends only on the parameter m/q . This implies that if we are playing a game with an opponent and wish to ruin him, then

1. if the game is in our favor (has a positive m) we wish to make m/q as large as possible, while
2. if the game is against us (has a negative m) we wish to make this quotient as close to 0 as possible.


This is intuitively plausible. It says that if the game is for us we wish to make the fluctuations small, and if the game is against us, we wish to make the fluctuations large. The above rule makes this obvious qualitative statement quantitative; it tells us (roughly) how to trade an increase or decrease in fluctuation for an increase or decrease in the average winnings or losses per play. The rule is, of course, quite approximate. It is easy to obtain better but less intuitive ones.

It is worth noting that the rule has consequences which are not widely known. For example, if one wanted to maximize his probability of winning a fixed amount, say \$10,000 at Las Vegas when starting from a smaller amount, say \$1,000, and had to choose between playing dice (average loss about 1.5% per play) or roulette (average loss about 5.3% per play) the choice of game depends not on the average loss but on the m/q . For typical Las Vegas betting limits (\$500 and \$100 respectively) contrary to the popular (and sometimes expert) belief, both games give about the same probability of success. In Part One we made the remark (p. 72) that in warfare (where presumably the aim is to bankrupt the opponent) the poorer contender generally wants to increase the variance (fluctuation) while the richer one tries to decrease it. This rule may seem to be in contradiction to the rule implied in Equation (4) which is not dependent on relative resources.

(footnote continued on next page)

Noisy Duel

Let us consider now a different kind of game, a duel game. Assume that you and your opponent are going to conduct a duel in which you start out some long distance apart and then will walk toward each other. You each have one bullet in your gun and both of you are going to shoot to kill. The accuracies of each of you get higher as you get closer. Now, it is clear that if you shoot first and miss, he will then kill you, since he then will not shoot until he gets right on top of you. Also vice versa. However, it is also clear that in general he may not wish to wait until you fire because you may not miss. If we assume that the accuracy of each of you starts from zero and increases in a linear way then the situation is indicated by Chart 9.



(footnote continued from previous page)

There are two reasons for the paradox. Equation (4) assumes that m and q are constants independent of the resources of the players. In war-like situations this is rarely true; m usually tends to favor the richer player. Secondly, we have ignored the non-zero-sum character of war (if we measure payoffs in resources) which tends to discriminate against the poorer player.

Each player in this game has an infinite number of alternative pure strategies available rather than just two. (A pure strategy consists of planning to fire at a certain distance away, if the opponent has not yet fired. Since there are an infinite number of distances, there are an infinite number of pure strategies.)

Let us now look at the graph in Chart 9. One of the lines gives the probability that you will kill him if you fire first. The other gives the probability that he misses you when he fires first. In other words, this line is given by one minus the probability that he kills you. Since you will then presumably wait until your accuracy is perfect before firing, it is also the probability that you will kill him when he fires first. We have marked in especially heavily the line which corresponds to the worst that can happen to you (the line of minimums).

As before the best thing that you can do is to take the maximum of this worst line, the so-called max-min point; this is at the 50% point. Therefore, if you have not yet walked 50% of the distance you will want to hold your fire and so will he. On the contrary both of you will desire to fire at the 50% point because if you don't the other guy may hold his fire a little while and then fire. He would then get more than one-half probability which he is entitled to. There is a question of what happens when you both fire at exactly the same instant but this happens to be trivial.⁷

⁷If the payoffs under different outcomes are given by the matrix
(footnote continued on next page)

In this case the optimal strategy was a pure strategy. Unlike some of the previous cases where we used a mixed strategy, neither player should pick a menu of possible actions mixed according to some probabilities. This occurs because the max-min of the pure strategies is equal to the min-max. For a similar game where it is best to use a mixed strategy let us consider a variation of this duel.

Silent Duel

The situation changes completely if you can't tell if the other man has fired or not. Under these circumstances, even though he has fired, it is not safe for you to wait until you are on top of him before firing, because you don't know that he has fired. It is possible that all the time that you are waiting he is also waiting and that his accuracy is increasing. Therefore you can't afford to wait too long even though you may think he has fired.

Let us first demonstrate that the previous pure strategy solution which was to fire at the 50% point is no longer "optimal." Assume, for example, you fire when the probability of killing him is $2/5$. Your opponent, however, does not know that you fired, and still uses the old strategy. He therefore fires at the one-half point.

Under these circumstances you have two chances in five of killing him

		He	
		survives	killed
You	survive	0	1
	killed	-1	0

then the payoff is continuous at the point of simultaneous firing; i.e., both players at that point can expect close to the value they would have gotten close to the point, in this case, 0. However if the numbers in the matrix are relatively different than the above, for example, one player doesn't much want to live anyway, then the game may fail to give such a precise solution to both players. One may be forced to always shoot a bullet's flight before the critical distance.

outright. His chance of killing you is given by the three chances in five that you have of missing him times the one chance in two he has of killing you. This comes to $3/10$. Lastly, there are two chances in ten that neither of you will be killed and will presumably fight the duel over again. Therefore, your relative chance of killing him is given by the ratio of $2/5$ or $3/10$ or 4 to 3. Since the game is symmetrical, there is no reason that you should have an advantage and it must be a mistake for him to always shoot at the 50% point. Actually you should both use what we have called a mixed strategy; that is, neither of you should decide on a definite place to fire but rather you should each choose and use a probability distribution that gives the probability that one will fire at a particular instant of time. Using a probability distribution makes it impossible for the enemy to know exactly what you are going to do.

The situation is formally similar to the coin matching game where we chose at random between heads and tails, only now we have an infinite number of choices available so we must use a probability density function instead of discrete probabilities. Chart 10 (the eventual solution) shows how in the optimal mixed strategy the probability of your shooting at any particular point should vary. You will note it has a slightly curious shape; as long as you have walked less than $1/3$ of the distance, you should never fire. Then your probability of firing should go up suddenly to a relatively high value and then gradually decrease. Roughly what happens is that you shouldn't fire until your accuracy gets to be at least appreciable and then you want to usually fire early rather than late; you want to get your shot in first. You do have to fire late at least once in a while to convince your opponent that it is unsafe for him to hold his fire. That is, you have to persuade him that it is important for him also to get a shot in early. If he happens

OPTIMAL STRATEGY FOR SILENT DUEL⁸




Chart 10

to know that you will never wait past some point, then if he happens to wait past that point, he can safely wait until he is on top of you. Also, it is important to fire late sometimes in order to take full advantage of the possibility that he may have fired and missed.

It is important to notice that even though this game is presumably played only once or at most a very few times, we still talk about the relative frequency of different choices. If this confuses the reader he can think of it in the following way.

Let him suppose that he is teaching military doctrine at a war college and he wishes to teach his students how to play this game but the advice that he gives will be made available to the opposition. He will then find that it is absolutely necessary to introduce mixed strategies even if each student is to play the game only once if--on the average--the enemy

⁸This distribution is given by

$$f(x) = \begin{cases} 0 & 0 \leq x \leq \frac{1}{3} \\ \frac{1}{4x^3} & \frac{1}{3} \leq x \leq 1 \end{cases}$$

students are not to get an unreasonable advantage. It is clear, therefore, that when one is giving advice, one may have to suggest mixed strategies. (It is also clear to us that when one is taking advice, even for "one-time" cases, he wants to use mixed strategies that maximize the probability of achieving his goals, but the justification of this course of action is a little harder to make convincing to some people and is not essential to our argument.)

Some Definitions and Formal Results

It may be well here to review some of the terminology we have introduced. We will give more attention to this than is necessary for this chapter because we would like to give the reader some of the flavor of the mathematics in the subject.

In order to have a game, one must have two or more players. (The reader should not be bemused by the possibility of solitaire; there is no way to get rich playing by himself.) There must be a conflict of interest among the players. That is, the game has several possible outcomes and the players have personal preferences for these, known to everyone. Also, each of the players have some control over the outcome. This control is naturally not complete. The other players may, when it is in their interest, be uncooperative, and there may be random chance events present. Lastly, there must be rules which each player knows and which give the complete range of alternatives available to each player. Game theory concerns the principles which guide intelligent action in these kinds of conflict situations.

There are several useful ways of classifying games. Games are either two-person, where there are only two sides with conflicting interest, or n person, where there are more than two sides. There is a rich theory for

the first of these, and a relatively sparse theory for the second. A game may have perfect information or it may not. Perfect information means that at every point in the game everything that has occurred in the game before that point is known to each of the players.

A game may be finite or infinite. It can be infinite in two different ways. First, there may be an infinite number of choices available at a move, such as in the duel games; secondly, the play may have the possibility of continuing indefinitely. Games commonly played for fun are generally finite.

Chess, and most parlor games, are described by the rules so that they consist of a series of moves made alternately by the players. A game described in that way is said to be in extensive form. Though this form makes the game more playable in the parlor, it is not necessary. A player may, rather than waiting to see what his opponent will do, simply describe what he will do in every contingency. This may be done for the next move, for several moves at once, or for the whole game. If he goes all the way, he has, in some sense, predigested the entire game beforehand. The list of the possible actions may be enormously long, but, at least in theory, it exists. Then a move for a player might consist of choosing such a possible list. If each player is to do this, and the moves are to be made simultaneously, the game is said to be described in normal form. (The reader should note that the word "normal" here carries the connotation of "standardized" rather than "typical.") It is in the "look before you leap" form as opposed to the original "don't cross your bridges until you come to them" form. A matrix game like matching pennies is, by its nature, in normal form. It is always possible to rewrite the rules of an extensive game so

that it is in normal form.

A strategy for a player is a consistent, determinate procedure for his actions during the play of a game. A pure strategy consists of a complete set of choices among the alternatives open to him, one for every situation which could arise. It is the choice of one move when the game is in the normal form mentioned in the previous paragraph. There is a further kind of strategy which is very important. A mixed strategy for a player consists of a set of probabilities or weights, i.e., a probability distribution, according to which he will choose pure strategies by chance. Such a strategy may guarantee a better average outcome for a player than any pure strategy. If the game is only available in the extensive (in the "don't cross bridges until ...") form, then one must restrict his strategies to a set of probability distributions on the alternatives available at each move as they arise in the game. Such a set is called a behavior strategy. Every behavior strategy is a mixed strategy but not vice versa. In many games, behavior strategies are as effective as mixed strategies. However, in other games this is not so. Behavior strategies can turn out to be not as effective as the mixing of pure strategies beforehand, and examples when this is so are easy to concoct.

The choice by each player of a strategy constitutes a play of the game. The result of a play is called the payoff of the game. In almost all formal game theory, it is supposed that the payoff represents a gain (possibly zero or negative) of an objective, transferable and numerical utility to each of the players. (Money is the best example of a transferable utility and for most purposes it can be considered as objective.) We can therefore think of the payoff as a function of the game strategies, the number of variables being the number of players.

The object of the game for each player is to maximize his expected utility. If the sum of the payoffs to the players in a game is always 0, we say it is a zero-sum game. In the case of a two-person zero-sum finite game, this function can be expressed as a payoff matrix, the elements being the amounts won by one player from the other. The first is called the maximizing player and the second is called the minimizing player.

If the rules are changed so that the maximizing player can choose after he knows his opponent has chosen, the new game is called the majorant game. If the minimizing player can choose second, it is called the minorant game. In these games, the best either player can do is use a pure strategy. Since in the majorant game the maximizing player will always choose a simple maximum, the minimizing player can do no better than pick a strategy which will minimize this maximum. The players' pair of strategies that does this is called the min-max point. The payoff at the min-max point, both intuitively and formally, must be at least as great as at the max-min point, and is generally greater.

The majorant and minorant games provide upper and lower bounds on what the two players could expect to achieve in the original game. It may be that these two bounds are equal. If so, then both players can achieve the common bound by resorting to pure strategies and the matrix of payoffs is said to have a saddlepoint, the element of the matrix where the min-max and the max-min coincide. This payoff is called the value of the game. It is the amount both players, in such a game, can guarantee themselves to achieve. Now it is an interesting fact the matrix of payoffs of any game of perfect information when put in normal form, possesses a saddlepoint, and it is therefore clear that both players need in this case use only pure strategies.

The central problem of game theory is to find the "best" strategy for each player in a game. Where there is a saddlepoint it is intuitively clear

that the corresponding strategies are the "best" strategies. There is a fundamental theorem for two-person zero-sum finite games which generalizes this result to games without saddlepoints. It states that in the domain of all mixed strategies, there are strategies for each player such that the min-max of the expected payoff is equal to the max-min. Using these strategies, a value of the game is determined between the values of the majorant and minorant games and a player can do no better on the average against a conservative opponent. Thus a saddlepoint is again achieved in the larger domain of mixed strategies.

The fundamental theorem actually holds for a much wider class than finite games. From a practical point of view, the principle embodied in the theorem is in some sense extendable to all reasonable two-person zero-sum games, though mathematicians interested in rigor discuss games in which there is no value in the simple sense above.

Attacking Targets of Unequal Importance

It may be interesting to indicate how complicated Part One of this book could have been if we had tried to introduce some game-theoretic arguments. For instance, let us assume as in the example of Part One that we have two airfields; but that one of the airfields has $2/3$ of our planes on it and the other airfield just has $1/3$, rather than each airfield having $1/2$ the planes. Assume also that we have two ground-to-air missiles of 100% accuracy and reliability so that each is guaranteed to shoot down an attacking plane. Finally assume that the enemy has two planes with which he intends to attack us and that he knows which field contains $2/3$ of our planes and which field contains $1/3$ but does not know how we deploy our missiles. We now have three possible choices.

1. We can put both of our missiles on the more valuable field. We will call this the (2,0) choice.

2. We can defend each target equally strongly. We will call this the (1,1) choice.
3. We can put both of our missiles on the less valuable target, the (0,2) choice.

The enemy has a similar set of choices for allocating his planes which we will label in the same way. Lastly, we will assume that if he gets one plane through the defenses of any airfield, that that plane will totally destroy the airfield. Let us now look at the payoff in terms of the proportion of the air force that is saved.

If both we and the enemy use the same choice the enemy gets zero because no planes get through. If we know what the enemy is going to do, we can defend ourselves perfectly. The payoff of the majorant game is 1.

attacker's choice defender's choice		strategy			strategy strategy	
		1/7	2/7	4/7		
strategy		(2,0)	(1,1)	(0,2)	1/7	2/7
4/7	(2,0)	1	2/3	2/3	2/3	5/7
2/7	(1,1)	1/3	1	2/3	1/3	5/7
1/7	(0,2)	1/3	1/3	1	1/3	5/7
column maximum		1	1	1	max-min = 2/3	
column average		5/7	5/7	5/7	min-max = 1	

maximum row average = 5/7
 minimum column average = 5/7
 max-min = 2/3 < 5/7 < 1 = min-max

Chart 11

If the enemy knows what we are going to do, he can always succeed in getting either $1/3$ or $2/3$ of our forces as shown by the fact that the row minimums vary between these two numbers. Under these circumstances we will use the max-min strategy of $(2,0)$ and the payoff of the minorant game is $2/3$. Since there is a difference between the max-min (minorant game) and the min-max (majorant game), at least one of the players should use a mixed strategy. It is easy to verify that both should and that the optimal mixed strategy is for we the defenders to play $4/7$, $2/7$, and $1/7$ while the enemy should play $1/7$, $2/7$ and $1/7$. Under these circumstances we save on the average $5/7$ of our force (and as expected, $5/7$ is between $2/3$, the maximum of the row minimums, and 1 , the minimum of the column maximums).

We are sure that many people will be uncomfortable over the fact that $1/7$ of the time we should play the fairly silly-looking $(0,2)$ alternative, where we put all of our defense on the less important airfield. Let us see what happens if we arbitrarily throw this alternative out. The new payoff matrix is shown in Chart 12.

attacker's choice defender's choice strategy		strategy			row minimum
		0	0	1	
		(2,0)	(1,1)	(0,2)	
1	(2,0)	1	$2/3$	$2/3$	$2/3$
0	(1,1)	$1/3$	1	$2/3$	$1/3$
column maximum		1	1	$2/3$	

max-min = $2/3$
min-max = $2/3$

Chart 12

Since now the max-min equals the min-max, both sides can afford to play a pure strategy. The defender plays the (2,0) alternative and the attacker should play the (0,2). The defender's losses then go up from $2/7$ to $1/3$ which is appreciable, but not really large. We are better off if the attacker doesn't realize that we have thrown out the (0,2) alternative because he probably still feels constrained to play the old $1/7, 2/7, 4/7$ strategy. In fact, if he is not vigilant, he practically has to do this because if he really plays the (0,2) alternative and we catch on, we can ourselves play the (0,2) alternative and make the payoff 1 which is a spectacular improvement. It is, we think, now clear why the defender plays the (0,2) alternative occasionally. He does it to prevent the opponent from playing his own (0,2) alternative exclusively and thereby gaining.

However, if we had any feeling at all that the assumptions on which the game was based might be in error or that there was a possibility that our security was not good and the enemy could tell how we were defending, then it might be best to simply defend the more valuable field and forget about trying to gain a .05 extra survival by following a tricky strategy.

In the reduced game, the defender's strategy could have been to play the first row $1/2$ of the time and the second $1/2$ of the time, or any strategy between that and the pure (2,0) strategy--the value to him is not changed. He has that alternative. The attacker, however, does not have any alternatives if he is to choose an optimum strategy. He must play the strategy shown or, against good play, make less than the maximum possible.

From the game theory point of view, the defender is indifferent between his two alternatives. But he may have some extra or non-game-theoretic preferences. For instance, he may think the attacker might be lax and once in a while try the alternative (1,1). Then the $1/2$ - $1/2$ strategy gives him

an advantage that the other does not. Or maybe he would also like to be prepared to play the game where the attacker will not use his (0,2) alternative at all either. Such considerations, while important to the real world, are irrelevant to our current narrow formulation.

It is more important for the attacker to use a mixed strategy. If the defender knows what the attacker is going to do, he can defend himself perfectly and change the value of the game from .71 to 1. However, if for any reason the attacker does not like one of his choices and throws it out, he doesn't lose very much. A calculation shows his loss is .06 in all three cases. This last result may seem slightly paradoxical to the reader. In the optimum mixed strategy the attacker plays the alternative (2,0) one seventh of the time; he plays each of the other alternatives much more often. Yet it is as serious for him if he arbitrarily omits this rather infrequent alternative as if he omits the more frequent alternatives.

This illustrates a fairly general point. It is not necessarily the frequency with which an alternative is played that makes it important. The mere existence of a possibility of playing a certain alternative is often sufficient to force the enemy to expensive countermeasures. Once the enemy has taken these countermeasures, it may no longer pay to play the threatening alternative very often--only often enough to keep the enemy "honest."

The Trader and the Cannibal

Let us now consider a completely different kind of game. Imagine for example, that you are a trader and are visiting Koko, chief of the cannibal island's gourmet club. You are in the following delicate situation.

You are going to give him a present of some beads. He is going to give you a present of some coconuts. If he considers

his present more valuable than yours, he will be insulted and have you seasoned and cooked. If he feels that your present is equal in value to his he will do nothing. If he considers your present more valuable than his, he will feel that he has lost face and let you have an extra present, an evening with his wife (fat, greasy and amorous), about whom you could not care less. Your only objective is to trade beads for coconuts.

The first problem we have to consider is the relative value of things to Koko and the trader; that is, these two people evaluate lives, coconuts, beads and wives quite differently. In fact, the heart of the problem lies in the fact that Koko values beads more than coconuts and the trader values coconuts more than beads so that it is conceivable that they can come to an amicable and mutually profitable arrangement. However, if we allow them to value things differently it makes the problem very difficult. Being at this point dedicated mathematicians, we will ignore what is the essence of the problem and assume that beads and coconuts are of equal value to both Koko and the trader. Let us do more than this. Let us also assume that the trader's life and Koko's wife are each worth three coconuts. Incidentally, these rather drastic assumptions are not being made merely for pedagogical reasons. We cannot really treat the problem in a non-controversial way unless we make some assumptions of this general type. However, let us continue on our way.

Chart 14 gives the payoff matrix which indicates what Koko and the trader get under various conditions.

PAYOFF TO KOKO										row minimum	row average
strategy	number of beads	1/6	1/6	1/6	1/6	1/6	1/6	0			
strategy	number of coconuts	1	2	3	4	5	6	7			
1/6	1	0	-2	-1	0	1	2	3		-2	0
1/6	2	2	0	-2	-1	0	1	2		-2	0
1/6	3	1	2	0	-2	-1	0	1		-2	0
1/6	4	0	1	2	0	-2	-1	0		-2	0
1/6	5	-1	0	1	2	0	-2	-1		-2	0
1/6	6	-2	-1	0	1	2	0	-2		-2	0
0	7	-3	-2	-1	0	1	2	0		-3	-1/2

column maximum	2	2	2	2	2	2	3	
column average	0	0	0	0	0	0	1/2	

$$\text{max-min} = -2$$

$$\text{min-max} = 2$$

$$\text{maximum (row average)} = \text{minimum (column average)}$$

Chart 11

Let us verify a few entries. If the number of beads and coconuts are the same then Koko nets zero. If Koko puts up 3 coconuts and the trader 1 beads then Koko gets 1 on the trade but loses 3 on his wife so he nets -2. If the trader puts up only 1 bead and Koko puts up 3 coconuts, then Koko

loses 2 on the trade but he gets to eat the trader which is worth +3; so he gets a net of +1, etc.

Under these circumstances it is easy to verify that both Koko and the trader should play a mixed strategy, such that $1/6$ of the time each of them is willing to put up 1, 2, 3, 4, 5 or 6 objects. They should never give more than six though. The strategy is verified by the usual process of calculating the expected value of all the columns and rows and showing that the max of the row averages is equal to the min of the column averages. This is done on the chart.

Let us now assume that the rules have been changed and that Koko is going to visit the trader on his ship. His canoe is such that it can carry at most three coconuts. The trader therefore knows that even though Koko would like to play the numbers one through six uniformly, he can actually only trade one, two or three coconuts. Therefore, if the trader plans on giving Koko four beads, he will automatically have beat Koko as far as the gift exchange goes. Now, of course, if Koko knows that the trader is going to give him four beads, he will in turn give just one coconut, the smallest he can. If both acted this way, the trade will be even, and the trader will have gotten no advantage even though Koko is limited in the number of coconuts he can carry. It turns out, of course, that the optimal plan for the trader is to play a mixed strategy, but a slightly different one from what he did previously. The situation is described in Chart 15.

strategy number of coconuts		strategy number of beads							row average	
		0	1/3	0	2/3	0	0	0		
		1	2	3	4	5	6	7		
strategy for bidden	2/3	1	0	-2	-1	0	1	2	3	-2/3
	0	2	2	0	-2	-1	0	1	2	-2/3
	1/3	3	1	2	0	-2	-1	0	1	-2/3
		4	0	1	2	0	-2	-1	0	1/3
		5	-1	0	1	2	0	-2	-1	4/3
		6	-2	-1	0	1	2	0	-2	1/3
		7	-3	-2	-1	0	1	2	0	-2/3
column average		1/3	-2/3	-2/3	-2/3	1/3	4/3	7/3		

$$\text{max-min} = -2/3$$

$$\text{min-max} = -2/3$$

Chart 15

It is shown on the chart that if Koko gives only one coconut two-thirds of the time and one-third of the time all three of the coconuts, he can guarantee that, on the average, he will lose at most $2/3$ (minimum of the column averages). The trader should, one-third of the time, be giving two beads and two-thirds of the time four beads. This guarantees that the most that Koko will get is $-2/3$ (maximum of the relevant row averages). Since

these two are equal this is the best method of playing the game.

We should notice, however, a very important thing. The trader is leaning rather heavily on the fact that Koko can carry at most three coconuts. If, for example, Koko suddenly produces four coconuts he will win on the average $2/3$ and if he produces five coconuts he will win on the average $4/3$, and, in fact, it should be clear that if Koko takes the game seriously, he is likely to build himself a slightly bigger canoe so that he can carry along some extra coconuts. In the world of sports or in the parlor, the rules of the game are set by officials and game theory may then work quite well. In practically no real life situations do the rules of the game have this sacred character.⁹

⁹ There is a continuous analogue of this game that is interesting. Assume that we irreparably allocate x resources in a battle and our enemy z resources. If $x > z$, we win the battle which is worth A to us. If $x < z$ we lose the battle which is worth $-B$ to us. If we tie, we get the average $(A+B)/2$. Precisely our payoff is:

$$\begin{array}{ll} A - x + z & \text{if } x > z \\ -B - x + z & \text{if } x < z \\ \frac{A+B}{2} & \text{if } x = z \end{array}$$

The optimal strategy then turns out to be for both players to play uniformly between 0 and $A+B$ and the value is $(A+B)/2$. Even though winning the battle is never worth more than A and losing B , we still occasionally allocate more resources than either A or B . If one of the players is limited so that he cannot allocate more resources than an amount a , i.e.,

$$x \leq a < A+B$$

then the other player can always force a win by allocating slightly more than a and the other player might as well not allocate anything. This, however, is not the optimal strategy. Instead both players should still play with the old density $1/(A+B)$ (uniformly between 0 and a) $a/(A+B)$ of the time. With the probability $[1 - a/(A+B)]$ left over, x plays 0 and y plays a . The value of the game is

$$V = a - B - \frac{a^2}{2(A+B)}$$

(footnote continued on next page)

If $A = B = 1$

$$V = - (1 - \frac{a}{2})^2$$

The above game demonstrates that in general even though it is valuable to be able to force a win, one may still not be willing to exert himself every time, and contrariwise, even though one knows that the enemy can force a win, that it still pays to occasionally exert oneself.

II. TWO-PERSON NON-ZERO-SUM GAMES

Two-person non-zero-sum games are sometimes thought of as a special case of three-person zero-sum games where the excess value is absorbed by a third (inert) player. Still, they are a class of special interest because they have one dimension less of flexibility, the alternatives open to the third player. Despite this, there is considerable difficulty in common, and from now on the theory will necessarily involve some new extra-mathematical considerations.

The Trader and the Cannibal (Continued)

Let us go back to our trader and his cannibal friend and try to take explicit account of the fact that actually each evaluates lives, coconuts, beads and women differently. It is clear, as a matter of fact, that if the trader, for example, has 100 beads, which mean nothing to him, and are worth a great deal to Koko, and if the cannibal chief has 100 coconuts which are worth nothing to him and a good deal to the trader, then there are the elements of a good trade present. If in addition Koko does not really enjoy taking the trader's life then the trader can expect to get away with 100 coconuts, his life, and a very unpleasant experience with the chief's wife. The chief will have his 100 beads, his digestion unimpaired, and his wife uninsulted. He doesn't really care about his wife and if they could have been frank, the trader would have been even better off.

There are, of course, a lot of other possibilities. The trader may be generous and give Koko ten thousand beads. At a later date, he may then find that Koko has enough beads and values them very little. This presumably means that the trader will have trouble in arranging future trades. Or the trader might end up with 100 coconuts and Koko with only 25 beads or vice versa.

Any particular distribution of beads and coconuts (also lives and wives) is called an imputation. The point is that there are a tremendous number of imputations available and nobody can really predict what will happen. What is even more to the point, nobody can say what either of them should do without knowing a great deal about the personalities and histories of the individuals involved.

In other words, the two-person non-zero-sum game involves more than just mathematics. At a minimum, it may literally include the principles of economics, politics, sociology, psychology, salesmanship, history, etc., for its treatment. What this usually means is that it can't be treated (in the sense of the considered opinion or scientific fact as opposed to the intuitive judgment). This is even more true in the n-person games, which we consider later.

Let us ask ourselves what a fair arbitrator or judge would decide is a reasonable bargain for Koko and the trader to enter into. Well, he might say that the two players should somehow gain equally from the trade. In practice, this is a very difficult thing to make numerical because it is impossible to measure Koko's satisfaction in having beads against the pecuniary profit the trader will get from the coconuts. If, however, both individuals have a common currency, say dollars, which they can exchange for many things, then the situation is much simpler. We can now measure Koko and the trader's surplus value in this medium. (The alert reader may notice that this last statement is superficial--we are very close to begging the question.)

There is still a serious ambiguity as to what we should call the trader's value. Should it be the maximum he is willing to pay for the coconuts or the minimum price for which he can buy coconuts from somebody else, or something

in between? We have a similar problem with Koko. From the viewpoint of a fair (socialist?) arbitrator, the first may seem correct, but from the viewpoint of the economist or businessman it is only the second which is relevant. If there is no market place in which the trader can buy coconuts or Koko can buy beads, then we are reduced strictly to a problem in "justice." Let us assume by fiat that this means both should gain equally in dollars. Under this very special circumstance of both narrowly circumscribing the environment, assuming they have a common medium of value and assuming that both should gain equally from the trade, we can solve the problem completely. We can, of course, also solve the problem when there is a market price for one or the other. But all of these situations are so restrictive that the reader may be tempted to say that calling them solutions is a little misleading. He is probably right, but we can claim they are sometimes interesting.¹⁰

¹⁰What actually happens, of course, in the case of a situation such as Koko and the trader is that it becomes rapidly institutionalized; the people involved react in a stylized pattern of behavior.

One of the writers happened to observe a rather interesting example of this situation. It seems that there is a large coat store in New York City which has a rather exclusive clientele. They are usually left with a fair amount of stock which they would like to get rid of at the end of the year. However, they do not want to get rid of it through local outlets because they think it would be bad if their regular customers or friends of their customers knew that they could buy the coats at reduced prices at the end of the year. They do not have a large enough volume of coats and the output is so variable, that they have not established a regular means of selling this remnant stock. However, over the past few years a friend of the writers, whom we shall call Alex (because that's his name), has been buying coats from them and selling them in a distant city. (The coat business is run by two partners, Sam and Al.)

The economic situation seems to be that the coats cost about \$50 to \$100 apiece to manufacture. They are worthless to Sam and Al and maybe less than worthless if they dispose of them in any other market than the traditional one represented by Alex. Alex buys them in lieu of a cheaper kind of coat

(footnote continued on next page)

which he can buy for \$20, so there is a trading margin between zero and \$20. The interesting thing is the way the exact bargain is arrived at during the bargaining negotiations. The process runs something like the following.

Alex comes to New York on a trip. He drops in to see his old friends Sam and Al. Al is not there, but Sam is, and Sam says, "Hi, Alex. How are things?" Alex says, "Fine. I'm here on a business trip, buying stuff." Sam says, "That's tremendous. We happen to have a lot of coats we can let you have." Alex says, "Oh, no. No. I just dropped in to say hello. I wouldn't dream of trying to buy your coats. They are much too good for my customers." (All this in seeming ignorance of the fact that he has bought coats from Sam and Al for the last five years running.) Sam points out that Alex has bought coats from him before, and they can probably make a deal. Alex explains that he has not come around to buy coats at all, he has merely dropped by to say hello, that the coats are much too good for his customers and it is, therefore, pointless to talk about him buying Sam's coats. Sam points out that it doesn't make any difference how good the coats are as long as Alex can buy them at a low price. Alex doesn't seem to hear and repeats that the coats are too good for his store and that he simply can't handle such a high-priced item. Sam demurs that the merchandise is indeed high-priced but not to Alex. Alex demurs that the merchandise is good quality and must be high-priced, and repeats it is much too good for his store. Sam says that Alex doesn't understand; he is giving the merchandise away, but just as a formality, so that it won't look like charity, Alex should make him an offer. Alex says, "No, I can't make you an offer. It would be too low, and you would be insulted." Sam says no, he is impervious to insults, just make him an offer. Alex says that he values Sam's friendship much too much to make him the kind of insulting offer he would have to make. Sam screams. "Make me an offer!" Alex says, "O.K. You asked for it. I think I could afford to pay you \$5 a coat." Sam turns purple, red and green, and then launches into a half hour tirade and cries, "Look at the lining, look at the buttons, look at the sewing, look at the style. Are you crazy?....." Alex is sorry that he has brought Sam so close to apoplexy and conjectures that he had better be on his way. Sam says, "Just a minute, please. Let me call up Al and see what he has to say." It turns out, of course, that Al is shocked by the offer of his old friend. He is willing to make a gift of the coats as a present, but if Alex doesn't want to accept a present, the price is \$10.

It turns out that after something close to four to five hours of arguing, mutual admiration, and threats, that the price is arrived at. It is invariably in the range \$10 to \$12. The exact price depends on the staying powers of Alex, Sam and Al.

The thing, however, which struck the observer most forcibly is the following incident. One time when Alex was in New York on a rush trip, he decided that he didn't have enough time to go through this four hours of arguing. The author suggested that since the thing had been a sort of ritual, he could afford to short-stop it and simply walk in, explain that he didn't have time and ask why couldn't they arrive at a price of \$12 without arguing. Alex' reply was that if he tried this, he might end up with only 30 minutes of bargaining, but the price would be nearer \$20, and furthermore a bad precedent would be set for future years. As a result, no sale was consummated and Sam and Al probably burned their coats. Alex went without his bargain. All would have been agreed that this is a small cost to pay for the preservation of a valuable social institution.

III. N-PERSON GAMES

We will now consider games with more than two players. We will discuss five games:

- (1) The Princess and Her Three Suitors
- (2) The Bankruptcy Court
- (3) A Pure Coalition Game
- (4) The Community of Shangri-La
- (5) The Game of Deterrence

The first game will illustrate the complexity and somewhat paradoxical-seeming results that can occur in even a simple three-person game. For example, under one form of rules where none of the players are allowed to get together it turns out that the most skilled player gets the least benefit from playing the game. A limited resolution of the paradox is obtained when rules are relaxed so that the players can form coalitions.

The Bankruptcy Court game is supposed to illustrate that it is possible to discuss some n-person game situations pretty well if one has an outside adjudicator who has definite principles to guide his actions.

The Pure Coalition Game illustrates the most characteristic feature of n-person games--the tendency of players to form coalitions and the pressures to double cross and to triple cross each other. While the game seems, in the form we present it, to be very simple, it turns out that almost all (the exceptions are trivial) three-person zero-sum games can be reduced to this form by a mathematical transformation. Therefore, once we understand this particular zero-sum game we have understood all three-person zero-sum games no matter how complicated the rules seem to be.

The game played by the community of Shangri-La is supposed to indicate

how societies can sometimes only get rational results by setting up what at first sight seem to be irrational institutions. Finally the last game, the Game of Deterrence, is supposed to explore some aspects of the notion of deterrence in more detail than we have done in the first part of this book.

The Princess and The Three Suitors¹¹

Let us consider the following very simple three-person game. A, B and C have decided to court the lovely Princess D. Her father, E, is a grouchy cuss and has given the three suitors the following proposition. "I will sit the three of you around the table. I may or may not put a mark on each of your foreheads--oops, this is the wrong game.¹² I will give each of you a gun. You will draw cards to decide in what turn you will shoot at one another. Once having established the order in which you will shoot, you

¹¹We are indebted to Lloyd Shapley for suggesting this example. There is a discussion of it by Martin Shubik in Readings in Game Theory and Political Behavior, Doubleday Short Studies in Political Science, No. 9, 1954.

¹²The king was thinking of how he married off Princess D's older sister, Princess L. He told each suitor, "I may or may not make a mark on each of your foreheads. I will then sit you around a table. Any suitor who sees a mark on any other forehead is to raise his hand. As soon as one of you figures out if he has or hasn't a mark, he should report to me." There is much cogitation, and then one suitor shouts that he knows that he has a mark. How does he know?

(This reminds us of some other riddles.)

- b. If you have 12 pennies and know that one and only one is off-weight, determine with a scale balance in three independent weighings which it is and whether it is lighter or heavier.
- c. Before counting the pennies, A, B and C had engaged in a bunch of track and field events. They amassed points (for 1st, 2nd and 3rd places) as follows:

A	22 points
B	9 points
C	9 points

(footnote continued on next page)

will continue shooting until only one of you is left. He will get the princess.

Let us further assume that A has a probability of $1/2$ of hitting anything he shoots at, B has a probability of $1/3$ and C has a probability of $1/4$. Now let us notice the following interesting effects. First, if all three are alive and rational, A will shoot at B and B will shoot at A. Neither of them will shoot at C because C, being the poorest marksman, is the least dangerous opponent. Similarly C will never shoot at either A or B, because if he succeeds in killing one of these dangerous opponents, the other will immediately proceed to shoot back at him. That is, he will fire in the air and wait until B has killed A or A has killed B. He will then take his turn and shoot at the survivor. He is thus guaranteed to get a first shot, before becoming a target himself.

We can, therefore, break the problem up into two pieces, the duel between A and B and then another duel between C and the survivor. Let us start by considering the first duel. If A happens to be lucky and shoots before B, then it is easy to calculate that A's chances of survival are $3/4$ and B's are only $1/4$. If B shoots first, they each have a survival of $1/2$. C always shoots first when he is tangling with one of the survivors. If he happens to fight with A, his survival probability is $2/5$ and if he happens

(footnote continued from previous page)

B got first in javelin throwing. Who placed second in the 100-yd dash?

- d. A big Indian and a little Indian were standing on a hill. The big Indian said to the little Indian. "You are my son but I am not your father." How can this be? Anyway, then the little Indian said: "Stop being silly."

If you give up on any of the above, see Appendix to this chapter.

to be fighting with B, his survival probability is $1/2$. The information is summarized on Chart 16.

(A,B)	(B,A)	(C,A)	(C,B)
$P_A = 3/4$	$P_B = 1/2$	$P_C = 2/5$	$P_C = 1/2$
$P_B = 1/4$	$P_A = 1/2$	$P_A = 3/5$	$P_B = 1/2$

Chart 16

Chart 16 shows what happens if A shoots first. His chance of survival is given by his chance of beating B, which is $3/4$ times his chances of then beating C, which is $3/5$ or a net chance of $9/20$. B's chance of survival is calculated in a similar fashion and comes out $1/8$, and C's chances come out $17/40$. The corresponding probabilities for the situation where B fires before A, are also shown.

<u>A Goes First</u>		<u>B Goes First</u>	
$P_A = (3/4) \cdot (3/5) = 9/20$	$= 18/40$	$P_A = (1/2) \cdot (3/5) = 3/10$	$= 12/40$
$P_B = (1/4) \cdot (1/2) = 1/8$	$= 5/40$	$P_B = (1/2) \cdot (1/2) = 1/4$	$= 10/40$
$P_C = (3/4) \cdot (2/5) + (1/4) \cdot (1/2) = \frac{17}{40}$	$\frac{17}{40}$	$P_C = (1/2) \cdot (2/5) + (1/2) \cdot (1/2) = \frac{13}{40}$	$\frac{13}{40}$

Since the two orderings are equally likely, we should average the results.

We then get:

$$P_A = (1/2) \cdot (18/40) + (1/2) \cdot (12/40) = 6/16$$

$$P_B = (1/2) \cdot (5/40) + (1/2) \cdot (10/40) = 3/16$$

$$P_C = (1/2) \cdot (17/40) + (1/2) \cdot (13/40) = 7/16$$

One immediately notices that C has the highest probability of success in this game even though his marksmanship is the worst; A is the second most likely

candidate for Princess D, and B has practically no chance at all. This is indeed a curious result. In fact from A and B's point of view it is a little offensive.

However, as long as everybody acts as an individualist and tries to optimize his personal probability, then the weakest player has the highest probability of success. It is because of this somewhat paradoxical result that we included this example.

Can A and B do anything about the result? Well, one way for them to improve the situation is for them to cooperate as a unit against C. As soon as they dispose of him they can toss a coin all over again to decide who should go first in the second round of shots, with A versus B.

The bargaining, however, may get a little complicated. For example, A may notice that if he and B get together in this way, C will spend all of his time shooting at A. That is, even though the game is heavily weighted against him, C will feel that if he is lucky and happens to kill one of his opponents he prefers to kill the most skillful one. Under such circumstances it can turn out that A's chances of winning the game are less than B's and he may well be miffed. The bargain helps B a great deal more than A.

About the only thing A can do is to bargain a bit more closely. He should try to get an agreement that if they happen to kill C, and then toss to see who is to go first, that they should not toss with a $1/2-1/2$ probability but with a biased probability which will somehow make up for A's loss.

We should also notice that it is not completely clear that if a coalition is formed it will necessarily be between A and B. A and C or B and C could get together. This possibility occurs because while it is best for C if there are no coalitions at all he may still be willing to work with one of his two opponents in order to prevent this opponent from joining with the

third man.

It is clear from the discussion that this fairly simple three-person game has unexpected subtleties in it. In practice it would be even worse because there would be the questions: can these three people trust each other when they make their bargains? Also they could all be lying about their accuracy. Therefore, one would not really know each man's probability of winning.

Probably the best thing for all of them to do is forget about the Princess D, who has a hairlip anyway, and just go home. But we as mathematicians cannot take this easy course. We must try to bring some sense out of this chaos. Some sense can in fact be made but not very much. For example, the three of them might get together and simply assign probabilities P_A , P_B and P_C for each to win the game. They will then draw a random number which will determine which one is to win the princess, the others to commit suicide. These probabilities of course are to be assigned in a fair manner.

The above is all to the good, if one can decide what is a fair manner. In fact, there are some statements which we can make about this too, though they are in no sense ultra-convincing. One might, for example, argue that these P_A , P_B , and P_C should be simply proportional to $1/2$, $1/3$ and $1/4$, the a priori probabilities, but C is going to be pretty hard to convince.

There are several other more mathematical "solutions" of this game. None of them are completely satisfying to the intuition, but they are worth discussing. It would not be right to discuss them all here. The one we will

discussion is due to Lloyd Shapley.¹³ This particular solution can be thought of as arising in the following manner:

The players are ordered in a random and unbiased fashion. Each player is then given a payoff equal to the extra value that he brings to a coalition formed by him and all the players ahead of him. The solution corresponds to the expected amount he will get.¹⁴

This means that each player first looks to see what he could get if both players combined against him. Then he looks to see how much he could add to a coalition if he joined up with another player and the probability of that coalition, and lastly, he sees how much they could all get together.

For example, the value to A is

$$\phi_A = \frac{1}{3} v(A) + \frac{1}{6} [v(A,B) - v(B)] + \frac{1}{6} [v(A,C) - v(C)] + \frac{1}{3} [v(A,B,C) - v(B,C)]$$

where $v(\dots)$ is amount the coalition can compel and ϕ_A is the adjudicated "fair" value.

Under these circumstances three "fair" probabilities to be assigned to A, B and C are .457, .310, and .233 respectively. The reader may not be fascinated by these numbers, but that is the way they come out.¹⁵

¹³ We should probably point out that it can be shown that the Shapley solution is the only one which has all of the following properties:

- (1) It doesn't discriminate between players as individuals but considers only their roles.
- (2) The value of the sum of two independent games is the sum of the values of the separate games.
- (3) The sum of the values is equal to the maximum possible value obtainable from the game.

In some sense anything that tries to call itself a solution should have the above properties or it cannot be used as a basis for adjudication.

¹⁴ The value to a coalition is defined in terms of a two-person zero-sum game played by the coalition against all other players. If the game is not automatically zero-sum, then a new zero-sum game is defined by letting the second coalition's payoff be $-T$ where T is what the first coalition gets.

¹⁵ We also considered the game where the players did not fire in turn but rather as a result of uniform random selection at each shot. In that case, the probabilities in the Shapley solution shifted slightly from the most skilled to the least skilled players, and were .447, .315, and .238 respectively.

Further Definitions and Formal Results (n-person games)

We will be content here to extend the concepts of two-person games to their natural generalization.

Two-person zero-sum finite games possess at least one pair of strategies which insures each player the value of the game. In n-person zero-sum games, the situation is much more complicated. We can first distinguish what can be called non-cooperative games. In these games, coalitions (open or secret) and side payments are not allowed. If there is a set of strategies, one for each player, such that no player can improve his outcome by deviating from his strategy if the other players do not, the set of strategies is said to constitute an equilibrium point. It can be shown that such equilibrium points exist for every n-person, zero-sum, finite non-cooperative game. They are in general not unique and only in certain cases do they have a common value of the game to the players.

When the players are allowed to cooperate, other considerations must enter. These are as much conceptual and methodological as mathematical, and are somewhat more refined than we should include here.

The Bankruptcy Court

The second n-person game we will consider is relatively tractable. It is the problem of the bankruptcy court. Assume that, for example, we have four creditors of the Bountiful Gold Mine Corporation and that the corporation has gone into bankruptcy. We will then have a situation where there are a lot of people with both conflicting and common interests. For example, nobody wants the corporation to go into a foreclosure sale where it will be picked up by some junk dealer for the salvage value of its almost non-existent assets. We will start by assuming that legally any one of the

creditors who is not satisfied can force such a foreclosure sale, but that the value of the corporation as a going concern is much greater than its liquidation value.

We will also assume that the four creditors of this corporation are the three suitors and princess of our previous example, A, B, C and D. (In fact, the king turned out not to be a king at all, but a gold mine speculator. Guess who bought stock?) Anyway, A has \$2.00 coming to him, B, C and D each have \$1.00 coming. Assume also that if the corporation is liquidated in an orderly fashion, it is worth about \$2.50 but that if it is liquidated at a forced sale it is worth only about \$1.00.

In other words, if the liquidation is orderly, the creditors can expect to average something like 50 cents on the dollar but if it goes to a forced sale, they can only expect to average 20 cents on the dollar. Under these circumstances, it pays the creditors to get together and agree to an orderly liquidation of the assets. However, if one of the creditors is obstreperous, he can presumably force the other creditors to pay him a bonus. For example, Princess D, exercising her womanly prerogative, might simply say, "Unless I get paid off 100 cents on the dollar, I will insist on a forced sale. Now, if you give me my full \$1.00, you will then have \$1.50 left to divide up among the three of you, which is better than the alternative of \$1.00 for all four of us." In fact, she can do worse than this. She can ask for a bonus of, for example, 25 cents and insist on getting \$1.25. It is one of the duties of the bankruptcy court to prevent this kind of behavior if it can do so without committing lese majeste. That is, the bankruptcy court will not insist that the creditors get together. That is the creditors' business, but the court will insist that under any arrangement that is made

all the creditors in the same class are treated on the same footing. If you have this sort of impartial referee available, then it is clear that the game should come to a reasonable solution, and everybody gets 50 cents on the dollar. The point that we wish to make though is that it takes an impartial referee to do it, and that the referee has fairly clear principles. Between different classes of creditors where the court doesn't have such clear principles, there is a partial rule of the jungle.

If we did things a little differently and tried to apply the Stanley solution, then we would see that there is a total surplus value possible of \$1.50, and since every man must contribute his vote in order to make this total possible, the surplus should be divided equally. This solution offends our legal and moral sense because we believe that people should get paid off somehow proportional to the amount they have legally coming to them and not to how much they could compel in a sort of "dog eat dog" situation. It is interesting to see what would happen in the "dog eat dog" situation if the rules were changed.

Assume now that we don't have an impartial referee but the rule is that if a majority of creditors, as measured by the amounts owing to them, get together, that any method of liquidation they agree to will be followed. Under these circumstances, A has a tremendous advantage. If he can get any one of the others to go along with him, he can force any kind of liquidation payoff he desires. If either B, C or D wish to fight him, he has to get two other people to agree. Let us ask ourselves what would be a reasonable way for A, B, C and D to evaluate their chances. That is, they have these paper claims against the bankrupt corporation; should they value these paper claims at the 20 cents on a dollar which can be obtained through a forced sale or the 50 cents which is what a fair court should get them or what?

Let us look again at the Shapley solution.

If A operated by himself he could force a payment of at least 40 cents. B, C and D operating by themselves can only force payments of 20 cents apiece. If A and one of the latter three get together the coalition can force a payment of 40 cents plus 20 cents plus \$1.50 or \$2.10. If we line the players up in random fashion, A has two chances in four (second and third positions) of ending up in a position where he can compel this increase and thus should get a bonus of $1/2 \times \$1.50$ or 75 cents. So A should get a total of $75 + 40$ cents or \$1.15. There is \$1.35 left over which is presumably distributed evenly among the other three, so the value to them is 45 cents apiece. The reader can easily verify from first principles that this figure is correct.

It is clear that because he has a large extra amount of power, A can compel an exorbitant return. We would have a somewhat different situation when there are only three creditors and two of them had \$2.00 apiece coming to them, and one had only \$1.00. In this case the \$1.00 player can compel (relatively speaking) an exorbitant payment, and in fact is entitled under the Shapley solution to 70 instead of 50 cents.

The results are intuitively reasonable. If there is one big fish alone with a lot of smaller fish, he can sort of dominate the situation but if there are two equally matched fish striving for supremacy, then any smaller fish who can tip the scales can get a great deal of benefit from his strategic position.

In some real sense the bankruptcy courts actually do follow the Shapley solution as between classes of creditors but within a class they insist on what is called non-preferential treatment.

A Pure Coalition Game

Let us look at the simplest possible three-person game. It is played with three people, Tom, Dick and Harry.¹⁶ Any two of these people can get together in the evening and form a temporary or permanent coalition. This coalition can force the other person to deliver at noon the next day a dollar to them. Then that evening a new coalition is formed or the old one is reaffirmed, and so on.

Let us now consider how one week's play of this game might go. On Sunday, it being a day of grace, and since everybody sort of thinks it is a silly game anyway, they don't bother playing it; everybody gets zero.

Sunday night, however, Tom and Dick get together and say to each other, "We are good friends, we never did like Harry anyway. Let's gang up on him and stay ganged up." Harry then loses a dollar.

Tuesday the same thing happens. By this time Harry is quite annoyed. That night he sees Tom, who is a pretty weak character, and says, "If you and me get together, I'll let you keep 75 cents." Tom isn't that weak. Before the bargaining is finished it is agreed that he is to get 90 cents. Therefore, Wednesday Tom comes out ahead 90 cents, Dick loses \$1.00 and Harry is ahead 10 cents plus the dollar he had been losing.

Dick is quite annoyed, sees Tom that night and berates him, but Tom is unmoved, so Thursday's payoff is the same. By Thursday night, Dick is just purple at his old friend Tom and goes to see Harry with the following proposition. "This is costing me a dollar a day and I am stuck with it.

¹⁶The intelligent reader will undoubtedly pierce the disguise our old friends A, B, and C have assumed.

However, I am hanged if I will stand for Tom getting 20 cents. If you won't break your agreement for only \$1.00, will you break it for \$2.00? That is, I will continue losing \$1.00 but I will give it to you as a side payment. In addition you and I will gang up on Tom and make him pay you an additional dollar."

We could go on, but it is clear that the party is getting rough. We show below a concise history of the payoffs before the blood begins to flow.

Day \ Player	Tom	Dick	Harry
Sunday	0	0	0
Monday	+ .50	+ .50	-1.00
Tuesday	+ .50	+ .50	-1.00
Wednesday	.90	-1.00	+ .10
Thursday	.90	-1.00	+ .10
Friday	-1.00	-1.00	+2.00
Saturday	?	?	?

It is clear that the possible combinations for crossing, double-crossing, and triple-crossing are infinite in this game. No one will be able to analyze it without some understanding of the individuals involved. However, there are certain things that can be said.

For example, the Wednesday, Thursday, and Friday imputations are somehow "irrational," and as we shall suggest, a set of imputations, such as given below, are "rational."

Imputation \ Player	Tom	Dick	Harry
One	+ .50	+ .50	-1.00
Two	+ .50	-1.00	+ .50
Three	-1.00	+ .50	+ .50

Such a set of imputations is sometimes called a "solution" of the game. While there is some disagreement about whether this is a reasonable use of the word "solution," the set does possess the following three interesting

properties:

1. There is no particular reason why all three or even any two rational people would feel compelled to switch from one member of the set to the other.
2. There is a definite compulsion for rational people to switch from any other imputation to one of this set. For example, Dick and Harry should prefer and together can force imputation three to Wednesday's solution and Tom and Dick should prefer and together can force imputation one to Friday's solution.
3. If by any bargaining chicanery one person can achieve any greater advantage than in one of the imputations of this set, he will almost certainly lose everything subsequently. Wednesday's solution, for instance, is very unstable for Tom, because Dick and Harry should prefer and can force imputation three.

The Curious Community of Shangri-La

Let us consider another game of a slightly similar character to our coconut trader and his dyspeptic friend but with a different set of overtones. This game is played by the entire community of Shangri-La, a community which is completely isolated from the world. Nightly, every individual in Shangri-La goes to the local temple and deposits a sealed self-addressed envelope which contains a sum of money, known only to him. The priests first sort these envelopes at random and match them up in pairs (there are an even number of people). They are then opened by an outside philanthropist who takes out (and keeps) the money in both envelopes and puts \$1.00 back in the envelope of each pair which had the larger sum. In case of ties, he tosses a coin to

see who gets the \$1.00. The envelopes are then collected by the priests and returned to the original owners.

We have asked many people how they would play this game. For some reason, most people say they would put in 50 cents or 99 cents. From society's point of view, 99 cents is obviously bad. If everybody did that there would be a net loss to the players and people would probably refuse to play. 50 cents is not much better. It means an average income of zero to the community.

In order to make the problem a little more dramatic, let us assume that the philanthropist is in fact the only source of income for the community and furthermore that the standard of living is such that it takes an average of 25 cents a day to survive. Therefore if everybody bets zero, and got, on the average, 50 cents a day income, not only would the community as a whole live very well, but even with fluctuation phenomena almost every individual would get a survival income.¹⁷

It is clear that it would be very reasonable for the authorities to compel everyone to make a zero bid. This maximizes the total income to the community, this total income is enough to support all reasonably, and given the distributive mechanism, it will be reasonably well allocated. The trouble is that there is no direct way under the rules of the game of finding out what any particular individual has bid. There is, therefore, no simple way to enforce such a rule.

We think it is clear how one might go about doing it in practice. One

¹⁷ The probability that under these circumstances any one player would average less than 25 cents a day over a year period is around 10^{-20} which is presumably small enough to be ignored.

would try to create a most drastic and violent theocracy. People would be raised from infancy up to believe that the one unforgivable sin is to put money in the envelope. It would, of course, still be true that some individuals would, either under the stress of some desperate temporary circumstances or because their religious training did not take, would, at least occasionally, put in a penny or two. Such individuals might eventually acquire great wealth.

There are at least three ways to handle this situation:

1. The community can simply ignore the chisellers and hope that so few people will cheat that it is not serious. They run the risk, of course, of a complete breakdown of morale and consequent disaster.
2. They could automatically shoot anyone who amassed a sum over some preassigned amount. While they would get some unfortunate people who had simply been lucky, they could set the limitations on what is an illegal wealth high enough to make this kind of mistake as infrequent as they wanted or could risk from the morale or distribution points of view. Illegal players would, of course, then hide their increased wealth. This would probably automatically reduce the temptation to get it.
3. A third thing they can do is to single out the wealthy as being a special group approved by the supernatural authorities. There would then be the natural inference that the others who are poor, are so because they are being punished by the same authorities for just thinking of cheating. (In this society everybody will have guilt feelings.) By thus holding up the wealthy as an example of rectitude, they can hope to reinforce the moral sanctions.

The Games of Deterrence--First Deterrence Game

We will now consider our last series of games,¹⁸ the games of deterrence. The first one is very simple. You and your enemy will be locked in a room. You both have a push button and the push button is attached to a keg of dynamite underneath the room.

¹⁸In more ways than one.

Your enemy says to you, "I want your wife and your fortune, otherwise I will blow both of us to kingdom come."

You say, "I don't believe it. You don't scare me one bit."

He says, "I really mean it." There is a gleam in his eye when he says it and you collapse. He walks off with wife and money.

Now, what is the trick in playing a game like this?


Well, the obvious thing is not to get into this game. Either make friends with your potential enemy or if you find you can't do that, spend what money you have in taking yourself out of what is practically an intolerable position. However, let us assume for the moment that neither of these two alternatives is available and that you really have to play.

You might as well admit from the start, that if you are playing this game with a madman you are going to end up minus wife and money or minus your life. Under these circumstances just make your choice. However, being a careful and judicious individual you have picked yourself a careful and judicious opponent. Given this, the game can be played in a reasonable fashion.

You have two choices. First, if your opponent is rational you might try the madman role yourself. (The reader should note that there is a very real payoff to making your announcement early). In this case you probably get his wife and money. However, bigamy is unlawful so you decide to act more reasonably. You would then commit yourself irrevocably to a contingent mutual suicide if the other guy steps too far out of line. He, being calm, reasonable, and judicious, also commits himself to a contingent mutual suicide if you step out of line.

It is clear that such mechanisms as anger, integrity, honor and public avowals will all be useful in this process of mutual commitment. The net

result if both of you really believe that the other is willing to commit suicide is that you will both live out your lives happily and peacefully with only a slight twitch and regular fees to a psychoanalyst.



Second Deterrence Game

The First Deterrence Game is really not so much a deterrence as a pure blackmail situation. The second game we are going to consider, while almost the exact opposite of the above game, is also not really what we would call Deterrence. It would be played as follows:

You have wired your opponent's house so that at any time you choose you could blow him up. He has wired your house in a corresponding fashion. Unlike the first game, this game pays off a tremendous premium to the man who goes first, rather than to the one who announces first. There can be two elements of stability present. First, you may not really be sure of the technical facts. For example, maybe the other guy has cut the wires or in some way tampered with your plan to blow him up or maybe his house is stronger than you expected or your dynamite weaker (i.e., you don't really trust your calculations). The other element of stability is the moral one. Murder is

forbidden. We put the two caveats in this order because between nations many moral questions seem to be less important than the uncertainty in the calculations. In any case, even with the caveats the situation is not very stable and our guess is that this game does not have a future.

Third Deterrence Game

The real Game of Deterrence is between the first two games and is played according to the following rules.

You have a reinforced concrete cellar in which a member of your family sits day and night. His job is to press the button that blows up your neighbor's house whenever the time seems correct. Your neighbor is similarly situated. The cellar may or may not be big enough to hold the entire family. Even if it is big enough to hold the entire family, one can't or doesn't want to live in the cellar 24 hours a day, so that one would need some kind of warning to save his family. Lastly, and very importantly, while each has calculated that his cellar will and his house will not withstand the enemy's dynamite, there is some uncertainty in the calculation.¹⁹ Neither is really

¹⁹ It is important to realize that mutual deterrence does not come about automatically because of the existence of dynamite, houses, and cellars, but:

1. only if all parties believe that the dynamite is strong enough to blow up the house but not strong enough to blow up the cellar,
2. both parties value their houses enough so that they are, in fact, restrained by the thought it may be destroyed or severely damaged.

It takes real work and vigilance on the part of both parties to maintain this situation. Therefore, a common statement, "Once we have a deterrent force of a certain size, more deterrence is unnecessary," may be untrue. It depends on the strength of your dynamite versus his house, and more importantly on his dynamite versus your cellar. Since these relationships depend at least partly on what the enemy does, he can raise or lower the ante required for deterrence.

The large uncertainty in the calculations have another important effect. Many people have noticed that all-out war in the twentieth century is probably

(footnote continued on next page)

sure that either or both of the cellars actually will take the dynamite blast. (The situation where the calculation runs the other way and only uncertainty stays the blow is too terrible to discuss again. It sort of reduces to the previous game.) Both, then, are willing to go to a lot of trouble to increase the force of his own dynamite and to strengthen his own cellar.

Now there are several points which can be made about this game. First, if you are acting very politely with your neighbor, then you will not need such a strong cellar because there is not very much reason for him to take the risk involved in trying to blow you up. (This is true even if he thinks your cellar might go with your house.) If, however, you are pushing him around or making life miserable for him (maybe by just frustrating his unreasonable desires) then your cellar had just better be very very good or he may take a chance and push his button. The main thing is that you can't use the threat of blowing him up for minor policing actions. For example, if your neighbor's boy steals apples from your apple tree, or if your neighbor's dog barks at night, it is kind of pointless to try to prevent this behavior by threatening to push the button. It is also pointless to make the threat, even if you think that the neighbor has egged on both the boy and the dog.

(footnote continued from previous page)


completely unreasonable to all participants. There seems to be two rational alternatives:

1. to reorganize the world so that large wars, if not all wars, are unnecessary,
2. to reform the institution of war itself, either by changing the technology or by both sides agreeing to limitations.


As far as large thermonuclear wars involving the participants' heartlands are concerned, almost all proposals on this last point that the writers have seen tend to be not only politically and socially unfeasible, but also conceptionally wrong because they ignore the often dominating effects of uncertainty.

You both realize that boys and dogs will inevitably do things which aggravate so there is no point in adopting a policy which inevitably will result in buttons being pushed at some fortuitous moment. Even if one wants to push the button it is better to pick the moment himself. Therefore, if one makes up his mind to use the dynamite as a reaction against minor irritations and these minor irritations are sure to occur, then he had better start shopping for tents. Even if the cellar is big enough to hold his entire family and strong enough to take the dynamite, he should still try to save the dynamite for serious affairs (but not necessarily as serious as when the cellar won't hold the whole family). It just is not worthwhile to have one's house blown up over a relatively minor and inevitable matter like a dog barking. Your neighbor can, in fact, feel so sure of this that if he is mean or nasty, he can afford to egg on both boy and dog.

Third deference game



Fourth deference game



Fourth Deterrence Game

This game is enlarged into an n-person game as follows:

There are other people on the block who may feel inclined to take sides with either you or your neighbor. Even more than that, they may be induced to come up with money with which one can buy more dynamite or better cellars.

However, there is a little gimmick in the rules which annoys these other people and makes them cautious. All of their houses are wired for dynamite so that either of the two main contenders can blow them up, either selectively or collectively. But the situation is not symmetrical. The "neutrals" have neither buttons nor cellars.

Being a third party on this block is kind of uncomfortable. A real estate agent would undoubtedly have a great deal of difficulty in selling one of these homes. But these people are stuck. They happen to live on the block, and transportation elsewhere is not available. Probably their reaction will be to try to ignore the whole situation, and being human, they will probably become really annoyed at anybody who brings up the precariousness of their position.

The interesting aspect of this fourth game is that there is now an extra value to both of the main opponents of having good strong cellars that will contain the entire family. If they don't have this kind of cellar then either one of these opponents can make all kinds of extreme threats toward the third parties and possibly succeed in forcing them to add their resources to his own. The other opponent may not be able to do much about it, except to emulate his opponent's behavior. If he tries any corrective kind of action, his family would be destroyed even if he personally survived in the safety of his cellar. If, however, he has a cellar which will contain in comparative safety all the things he holds precious, he can (but is not

likely to because he is still unwilling to sacrifice his house) present his opponent with an ultimatum if his opponent really indulges in very reckless or provocative behavior. If the cellar is appropriate, it is almost impossible for the opponent to counter this strategy. Even preempting won't help the opponent because even if he preempts, and destroys the other's house plus family, he is still guaranteed to lose his own house and has, therefore, just won a Pyrric victory. Contrariwise, if he waits for an ultimatum to be delivered, he can be sure that the person who makes the ultimatum has already put his family in a place of safety. If he delivers an ultimatum of his own, the recipient is then warned and again is sure to put his family in the cellar.

As we mentioned, while it is true that neither of the opponents is likely to deliver an ultimatum lightly because even if one can save his immediate family he would still lose his house, the existence of the cellar makes the delivery of an ultimatum credible. Once both sides find ultimatums credible then they may be deterred from certain kinds of provocative behavior toward the "neutrals" as well as to each other. If they aren't, then you have a real problem.

This raises the interesting question of what kind of things one can expect to deter. It is clear to the writers that the time sequence may be all important here. If one sees his neighbor digging up his apple tree, he may be just mad enough to blow him up even though it doesn't pay to trade his house and risk annihilation just for the sake of an apple tree. Because your enemy knows that there is a strong possibility that you will act irrationally, he will probably be deterred from such a flagrant violation of the peace of the neighborhood. However, if he can depend on you thinking about it before you acted (if, for example, the power was going to be turned

off for 24 hours or if you had to call a family conference) then he might take a chance. Insofar as deterrence depends on one acting irrationally in cold blood, it is a weak defense. In particular, the complete or partial mutual suicide threat has only a limited usefulness in deterring limited provocation²⁰ that can be made to look ambiguous or which can be presented in a context that forces the opponent to carefully consider the consequences.

Fifth Deterrence Game



The fifth deterrence game is exactly the same as the fourth deterrence game with the addition of a research and development program. We assume now that both players are trying to develop better bombs and better concrete for their cellars. It is clear that if one player gets a substantial lead on the other player, so that for example his bomb is certain to wreck his enemy's cellar making it impossible for the other to retaliate, then the quality of the game will

change drastically. Under these circumstances it behooves both players to have extremely large research and development programs and to follow up all

²⁰The firm committal to take unlimited and in effect self destructive measures in order to deter important but limited provocations is sometimes called the "rationality of irrationality." The same reaction to unimportant or very limited provocations might be called the "irrationality of irrationality."

P-1166
7-30-57
- 64 -

the interesting possibilities that they can afford to. All of the considerations we mentioned in Part One about the nature of the decision process now become relevant.

APPENDIX

Answers to Riddles (p. 11)

- a. With that much cogitation, it is the only possible situation.
- b. Try (1)(2)(3)(4) ~ (5)(6)(7)(8), (1)(2)(9)(10) ~ (3)(6)(10)(11), and
(1)(4)(7)(10) ~ (2)(3)(11)(12).
- c. This is almost impossible to explain but any reasonably bright person can figure out it has to be C.
- d. She's his mother.

Bibliography

- 1. Von Neumann and Morgenstern, Theory of Games and Economic Behavior, Princeton, 1944.
- 2. Dresher, M., Theory and Applications of Games of Strategy, R-216, 1951.
- 3. McKinsey, J., Introduction to the Theory of Games, McGraw-Hill, 1954.
- 4. Williams, John, The Compleat Strategyst, McGraw-Hill, 1954.
- 5. Luce and Raiffa, Games and Decisions, John Wiley and Sons, 1957.
- 6. Shapley, L. S., The Value of An n-Person Game, RM-670, 1951.
- 7. Kuhn, H. W., Lectures on the Theory of Games, Princeton, Spring Term, 1952.